Prof. Bryan Caplan
bcaplan@gmu.edu
http://www.bcaplan.com
Econ 812

Weeks 3-4: Intro to Game Theory

I. The Hard Case: When Strategy Matters
  A. You can go surprisingly far with general equilibrium theory, but ultimately many people find it unsatisfying. In the real world, people frequently stand in between the one-agent and the near-infinite-agent poles.
  B. Even when people start out in the near-infinite-agent case, they often ex post end up interacting with a few people.
    1. Ex: Marriage market
  C. Game theory tries to analyze situations where strategy does matter. It generally ends up with less determinate answers than GE, but is often arguably more realistic. ("I'd rather be vaguely right than clearly wrong.")

II. Extensive and Normal Forms
  A. Standard consumer choice provides the basic building blocks: game theory retains the standard assumption that people maximize utility functions. Slight change: Game theorists often talk about "payoffs" instead of utility. The concept is the same: Given a choice of payoffs, agents pick the largest.
    1. Payoffs are usually interpreted as von Neumann-Morgenstern utilities to sidestep issues of risk aversion.
  B. Any game can be represented in two different ways: *extensive* form and *normal* form.
  C. Extensive forms display every possible course of game events, turn by turn. They show how behavior branches out from "choice nodes," showing payoffs at the end of each branch as it ends. For this reason, extensive forms are often called "decision trees."
  D. Simple example: Your career game tree. At each node you can keep going to school, or get a job and get your payout.
  E. More interesting example: *The French Connection* subway game. Criminal decides whether to get on or off the subway; then Popeye decides whether to get on or off. From the first node, the tree spreads out into two branches; then each of those branches spreads out to two further branches; then the game ends. Payoffs for {Criminal, Popeye}: (on, on)=(0,10); (on, off)=(10,0); (off, on)=(10,0); (off,off)=(0,10).
  F. Complications:
    1. Nature as a random player.
    2. Information sets: simultaneous moves are equivalent to sequential moves with uncertainty.

3. If you learn something before you decide, node representing what is learned must precede node where decision is taken.

G. Normal forms (aka "strategic forms"), in contrast, display a complete grid of strategy profiles and payoffs. The grid has one dimension per player.

1. Important: Strategy profiles often contain irrelevant information about what you would have done in situations that did not in fact arise.

H. Normal form of your 1-player career game:

Drop out before H.S.	Finish H.S., stop	Finish B.A., stop	Finish Ph.D., stop	Finish 2 Ph.D.s, stop
10	15	20	30	0

I. Normal form of the *French Connection Game*:

	Popeye		
Criminal		On	Off
	On	0,10	10,0
	Off	10,0	0,10

J. Example from Kreps: Player 1 chooses A or D. If D, game ends. If A, then player 2 chooses α or δ. If δ, game ends. If α, player 1 chooses a or d, and either way, the game ends.

K. Normal form:

	α	δ
Aa	3,1	4,3
Ad	2,3	4,3
Da	1,2	1,2
Dd	1,2	1,2

L. Challenge: Write down the extensive form.

III. Strictly and Weakly Dominant Strategies

A. So what does game theory claim people do? It begins with some relatively weak assumptions, then gradually strengthens them until a plausible answer emerges.

B. Weakest assumption: People do not play strictly dominated strategies. If there is a strategy that is strictly worse for you *no matter what* your opponent does, you do not play it. If elimination of strictly dominated strategies leaves you with a single equilibrium, the game is *dominance solvable*.

C. Classic example: Prisoners' Dilemma.

D. If all players think this way, you can extend this idea to *successive strict dominance*. If your opponent would never play a strategy, you can cross out that row or column. This may in turn imply that some more of your strategies are strictly dominated, and so on.

1. Fun fact: Order of iteration does not matter.

E. A dominance solvable normal form from Kreps:

	t1	t2	t3
s1	4,3	2,7	0,4
s2	**5,5**	5,-1	-4,-2

F. Further refinement: If probabilistic combination of strategies strictly dominates another for any probability distribution, that too may be eliminated. Then this normal form from Kreps becomes dominance solvable:

	t1	t2	t3
s1	**4,10**	3,0	1,3
s2	0,0	2,10	10,3

G. It may happen that one strategy is sometimes strictly worse and never strictly better than another. Using the criterion of *weak dominance*, such strategies may also be eliminated. Unfortunately, with weak dominance, order of iteration may matter.

IV. Backwards Induction

A. In any game perfect information, each node marks the beginning of what can be seen as *another* game of perfect information.

B. Question: What happens if we apply the procedure of "backwards induction," i.e., repeatedly apply strict dominance to these "subgames"?

C. Intuition: Systematically reason "If we get to this point in the game, no one would even do such-and-such, so we can erase that part of the tree."

D. Modest Answer: We can eliminate more possibilities than before.

1. Consider extensive and normal forms from Kreps **(Figure 12.5)**.

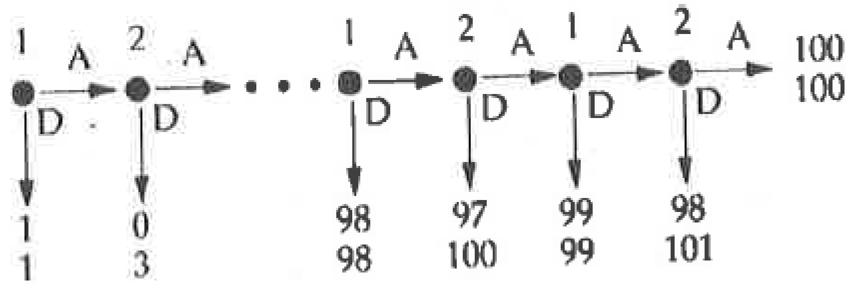

Figure 12.6. Rosenthal's centipede game.

E. Immodest Answer: Any finite game of complete and perfect information without ties becomes dominance solvable.

1. Chess example

F. Ex: The Centipede game **(Figure 12.6)**

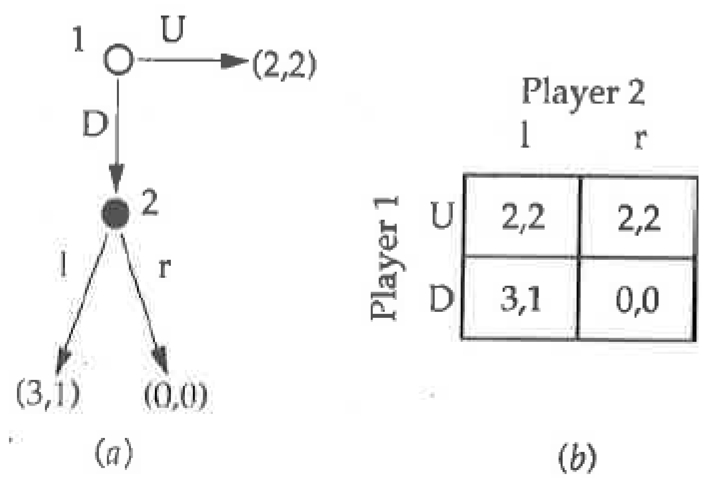

Figure 12.5. A simple extensive form game and its normal form counterpart.

V. Pure Strategy Nash Equilibrium

A. You can only get so far with strict dominance-type reasoning. Backwards induction seems impressive at first, but it only works for finite games of perfect and complete information. Very few interesting situations fit that description.

B. This leads us to a very different equilibrium concept, the *pure strategy Nash equilibrium*. A set of player strategies is a PSNE if and only if NO player could do *strictly* better by changing strategies, **holding all other players' strategies fixed**.

1. Imagine asking players *one-by-one* if they would like to do something different. If ALL of them answer **no**, you have a PSNE.
2. From the definition, it should be obvious that a game can have multiple PSNE or zero PSNE.

C. Example #1. Find the PSNE. How does this differ from strict dominance?

		Player 2	
Player 1		Left	Right
	Up	15,10	**8,15**
	Down	10,7	6,8

D. Example #2: Find the PSNE. How does this differ from strict dominance?

		Player 2	
Player 1		Left	Right
	Up	10,10	**0,15**
	Down	**15,0**	-5,-5

E. Example #3: Note the absence of any PSNE.

		Player 2	
Player 1		Left	Right
	Up	10,0	0,10
	Down	0,10	10,0

F. The PSNE concept is probably the most used in game theory and modern economics generally. It is somewhat paradoxical, however, because it seems to assume away strategic interaction, precisely what game theory was intended to address! A more strategic player might think "I'm not going to switch just because I would be better off holding my opponent's action constant. Maybe he'll *respond* in a way that makes me wish I hadn't changed in the first place."

VI. Mixed Strategy Nash Equilibrium

A. Talking about "pure strategy" NE strongly suggests a contrasting concept of "mixed strategy" NE. Instead of just asking whether any player has an incentive to change strategies, you could ask whether any player has an incentive to change his *probability* of playing various strategies.

B. How do you solve for MSNE? Each player has to play a mixture that leaves *all other* players indifferent.

C. Ex: Return to the game where:

		Player 2	
Player 1		Left	Right
	Up	8,10	1,15
	Down	12,0	-9,-5

D. When is player 2 indifferent between playing Left and playing Right? Let player 1's probability of playing Up be σ, and Down be $(1-\sigma)$. Then player 2 is indifferent so long as: $10\sigma + 0(1-\sigma) = 15\sigma - 5(1-\sigma)$, which simplifies to: σ=.5.

E. When is player 1 indifferent between playing Up and playing Down? Let player 2's probability of playing Left be φ, and Right be $(1-\varphi)$. Then player 1 is indifferent so long as: $8\varphi + 1(1-\varphi) = 12\varphi - 9(1-\varphi)$, which simplifies to φ=5/7.

F. So there is a MSNE of (σ,φ)=(.5, 5/7). When player 1 plays Up with probability .5, and player 2 players Left with probability 5/7, neither could do better by changing their mix. (They wouldn't do worse either, admittedly!).

G. Many people find the MSNE bizarre, but I maintain the opposite. The MSNE concept brilliantly accommodates the strategic complexity of real-world small-numbers interaction. Think of it this way: You make your opponents indifferent in order to *eliminate behavioral patterns they could exploit.*

1. Ex: Sports. You don't do the same thing all of the time because opponents will notice the pattern and play the most effective response. A predictable player is easy to beat. In

racquetball, for example, you play a mix of hard and soft serves, aiming at different locations on the court.

2. Ex: Strategy games. If you always attack the same place, your opponent will put all of his defensive strength there. In Diplomacy, for example, you randomize your attacks because a fully anticipated attack is easy to repel.
3. Ex: Rock, Paper, Scissors. You randomize to avoid being a sucker. Of course, if you play against someone who doesn't randomize, you don't want to randomize either; but maybe they are just tricking you into *thinking* they don't randomize!
4. Ex: Bargaining. If you are a hard bargainer, you get better but fewer deals. If you are a soft bargainer, you get worse but more deals. Which strategy works better? Neither!

H. MSNE cuts the Gordian knot of unlimited second-guessing, third-guessing, etc. All of these layers of thought can be reinterpreted as a randomizing device.

I. Solve the *French Connection* game. (Note the parallels to the Austrians' Sherlock Holmes example).

VII. Subgame Perfection

A. Suppose I threaten to fail any student who leaves early from any class. If you believe my threat, you will not leave early, and I will never have to impose my threat. This sounds like a Nash equilibrium - since I get what I want at no cost to me, and you prefer sitting in class to failing, neither wants to change.

B. But this sounds like an implausible prediction, because I probably would not want to carry out that threat. There would be a big fight, I would have to explain myself to the chairman, the dean, etc. How can a threat I would never carry out change your behavior?

C. In general terms, this is known as the problem of "out of equilibrium" play. I can optimally choose bizarre behavior in situations that I know will never happen. But knowing what I *would* do in situations that will never happen can affect your actual behavior in situations that routinely happen!

D. This gives rise to the Nash refinement of *subgame perfection*. Subgame perfection, in essence, requires Nash play in every subgame of a game.

E. To check for subgame perfection, you apply backwards induction as far as you are able. Thus in games of perfect and complete information, the result you get from backwards induction is always subgame perfect.

F. Standard example: Entry game. The two PSNE are (In, Accommodate) and (Out, Fight). But only the first is subgame perfect.

G. In games of imperfect information, though, you have to switch from strict dominance to Nash.

VIII. Prisoners' Dilemma

A. Surely the most analyzed game in economics is the Prisoners' Dilemma. Standard representation:

		Player 2	
Player 1		Coop	Don't
	Coop	5,5	0,6
	Don't	6,0	**1,1**

B. Natural solution concept: Strict dominance. Player 1 is better off not cooperating no matter what Player 2 does. Player 2 is better off not cooperating no matter what Player 1 does. So neither cooperates.

C. The Prisoners' Dilemma has many applications: public goods and externalities, collusion, voting, revolution... Others?

D. There is a lot of experimental literature on the PD. The extreme prediction is rarely borne out (people will cooperate even when defection is strictly dominant). But people do "leave money on the table," and there are a number of standard ways to reduce cooperation levels.

E. Moreover, no experiment that I know of has people play for, say, a year. I would strongly expect large-N, long-term play to closely match the game theoretic prediction.

IX. Coordination Games

A. Another game with a high profile in both theoretical and policy discussions is the Coordination game. Standard representation:

		Player 2	
Player 1		Left	Right
	Left	**3,3**	0,0
	Right	0,0	**5,5**

B. Natural solution concept: PSNE. If Player 1 plays Left, Player 2 is better off playing Left. If Player 1 plays Right, Player 2 is better off playing Right. And vice versa.

C. Coordination games underlie the whole path-dependence literature. Main idea: It is *possible* for people to be "locked-in" to Pareto inferior equilibria. (Of course, mere possibility is hardly proof!)

D. Problems like this naturally lead us to the notion of focal or "Schelling" points. Some coordination equilibrium are in some sense more obvious than others.

1. The classic NYC meeting example.

E. What would it take to actually get people into the Pareto-inferior NE? Most plausibly, at least a moderate number of players and gradual information dispersion.

F. Experimental evidence? Not too surprising.

X. Ultimatum Games

A. The Ultimatum Game is another game that has received a lot of academic attention. Standard set-up: Player 1 proposes one way to divide $10 between himself and Player 2. Player 2 accepts or rejects the division. If he accepts, they get Player 1's proposal; if he rejects, they both get 0.

		Player 2	
Player 1		Accept	Reject
	t	(10-*t*),*t*	0,0

B. Natural solution concept: Subgame perfection. Player 2 will accept any amount greater than 0, so Player 1 offers $.01 and takes $9.99 for himself.

C. Experimentally, no one does this. Even splits are common, and people often reject "ungenerous" offers.

D. Is this motivated purely by spite? Parallel Dictator game proves otherwise.